



Institute for Scientific Computing Research

ITS Lecture Series and Lecturer Biosketches



August 30, 2000

The Design of Design

Frederick P. Brooks, Jr.

University of North Carolina

Email: brooks@cs.unc.edu

Abstract:

The processes of designing computers, programming languages, operating systems, and big applications seem to have a lot in common. More interestingly, these same commonalities are also found in the processes of designing buildings, and even organizations. In some fields, such as building architecture and industrial design, there is currently a lot of study of the design process. Perhaps by studying the process, across the various media in which designs are realized, we can learn how to better practice, teach, and design.

Engineers have a simplistic, deterministic model of their own design processes. As soon as one expresses the model, real designers acknowledge it and recognize its gross inadequacies. Architects rejected it decades ago. What is a better model?

Through the centuries, designers have debated the fundamental issue of whether design should be rational or empirical, a question with many corollaries and implications. This talk fiercely defends the empirical side.

Great works of art have typically come from one mind; many great works of design have the same property. Now, when time-to-market pressures and increased specialization of skills dictate design teams, how shall we achieve conceptual integrity in our designs? What are the implications for telecollaboration?

Finally, I shall argue that "Great designs come from great designers, not from great product processes," and that many product troubles come from terribly wrong consensus-based design processes.

Speaker's web page: <http://www.cs.unc.edu/~brooks/>

Institution web page: <http://www.unc.edu/>

Biographical Sketch:

Frederick P. Brooks, Jr

Frederick P. Brooks, Jr., was born in 1931 in Durham, North Carolina. He received an A.B. summa cum laude in physics from Duke University and a Ph.D. in computer science from Harvard, under Howard Aiken, the inventor of the early Harvard computers. Dr. Brooks then joined IBM, working in Poughkeepsie and Yorktown, New York, from 1956 to 1965. He was an architect of the Stretch and Harvest computers and then was the project manager for the development of IBM's System/360 family of computers and Operating System/360 software. For this work he received a National Medal of Technology jointly with Bob O. Evans and Erich Bloch.

In 1957 Dr. Brooks and Dora Sweeney patented an interrupt system for the IBM Stretch computer that introduced most features of today's interrupt systems. His System/360 team first achieved strict compatibility—upward and downward—in a computer family. Brooks coined the term “computer architecture” in relation to the System/360 family. His early concern for word processing led to his selection of the 8-bit byte and the lowercase alphabet for the System/360; he engineered many new 8-bit input/output devices and provided a character-string datatype in PL/I.

In 1964 he founded the Computer Science Department at the University of North Carolina at Chapel Hill and chaired it for 20 years. Currently, he is Kenan Professor of Computer Science at UNC. His principal research is in real-time, three-dimensional, computer graphics “virtual reality.” His research has helped biochemists solve the structure of complex molecules and has enabled architects to “walk through” structures still being designed. He is pioneering the use of force display to supplement visual graphics.

Dr. Brooks distilled the successes and failures of the development of Operating System/360 in the book *The Mythical Man-Month: Essays in Software Engineering* (1975, 20th Anniversary Edition, 1995). He further examined software engineering in his well-known 1986 paper, “No Silver Bullet.” In 1997, he and Professor Gerrit Blaauw published a major research monograph, *Computer Architecture: Concepts and Evolution*.

Brooks has served on the National Science Board and the Defense Science Board. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences. He has received the ACM A.M. Turing Award, the IEEE John von Neumann Medal; the IEEE Computer Society's McDowell and Computer Pioneer Awards, the ACM Allen Newell and Distinguished Service Awards, the AFIPS Harry Goode Award, an honorary Doctor of Technical Science from ETH-Zurich, and The National Medal of Technology (from President Reagan) in 1985.

June 9, 2000

Mathematics and Computers

Peter Lax

Courant Institute of the Mathematical
Sciences, New York University

Email: lax@cims.nyu.edu

Abstract:

This talk will give an overview of the rise of computing in mathematics, and the infusion of mathematical ideas into computing in the past half century. The founding father of computational mathematics is von Neumann, who first set the agenda for the development of programmable electronic computers with large memory, as well as programming languages, schemes of discretising continuum equations, and algorithms for solving the discretised equations.

Everybody with the least familiarity with computing is aware of the continuing phenomenal improvement in the last fifty years of computer speed, memory capacity, software, and graphics capabilities. Problems that at one time have strained computing capacities could, a few years, later be done routinely, cheaply, and fast. But most people are unaware of how much this progress is due in equal measure to new mathematical ideas for tackling computational tasks. Striking examples include: shock capturing, multigrid, the fast Fourier transform, fast matrix multiplication, multiresolution, neural nets, etc. No doubt many more such mathematical shortcuts remain to be discovered.

von Neumann also realized that computing can do more than grind out by brute force numerical answers to concrete questions; it can discover new phenomena. Theorists can use computing in the manner of the experimentalists. In just such fashion did Fermi, Ulam, and Pasta discover recurrence in the dynamics of nonlinear chains, Kruskal and Zabusky solitons, Feigenbaum period doubling, Lorenz strange attractions, and Mandelbrot fractals. Perhaps the most astonishing discovery through numerical experimentation is Odlyzko's exploration of the connection between the distribution of the zeros of Riemann's zeta function and the distribution of the eigenvalues of random matrices. To prove rigorously what has been discovered by computation is a tremendous challenge of the future.

Computing has also made its mark as an aid in proofs. The most spectacular example is Haken and Appel's demonstration of the four-color theorem, where the computer is required to perform exact calculations. But computing is also an integral part of the logical structure of some analytical proofs, where the computations are only approximate. One can look forward to more and more computer assisted proofs. Have fun!

Institutional web page: <http://www.nyu.edu/>

Biographical Sketch:

Peter Lax

Peter Lax received his Ph.D. in 1949 from the Courant Institute of Mathematical Sciences at New York University. He then went to Los Alamos, where he had been stationed in the Army during the last year of World War II. He worked with John von Neumann, devising schemes for calculating flow with many interacting shocks. His first numerical experiments were carried out in 1952 on Los Alamos' MANIAC computer.

In 1954 he returned to the Courant Institute to join the Atomic Energy Commission Computing Center. He has spent the remainder of his professional career at NYU, making leading contributions to computational mathematics. His numerous honors include the National Medal of Science, the Lester R. Ford Award, the Chauvenet Prize, the Semmelweis Medal, the Wiener Prize, and the Wolf Prize. Many of his students have themselves become leaders in scientific computing.

Among his many contributions to the High Performance Computing community, one of the most important was his role in establishing the National Science Foundation (NSF) computing centers. He served on the National Science Board from 1980 to 1986 and chaired the "Lax Panel," which recommended that the NSF establish five national computing centers to provide university scientists access to the same level of high-performance computing available in the national laboratories. The articulation of this requirement in national science policy helped spur the birth of all subsequent HPCC programs, including ASCI.

May 12, 2000

How Shall We Program High Performance Computers?

Burton Smith

Tera Computer Company

Email: burton@tera.com

Abstract:

Uniprocessor computer architecture has traditionally been motivated by programming languages and operating systems, with benchmarks written in the usual languages also having some influence. In high performance computing the situation is curiously reversed, with architecture determining the principal characteristics of programming languages, operating systems, and benchmarks. The result has been chaos; a “software crisis” has been declared, and better tools for the development of parallel software have been demanded. The outlook for good tools is bleak without a new approach to the problem, which should include the engineering of computer systems with both system and application software in mind and the development of programming abstractions that are both effective and efficient on hardware we can build.

Speaker's web page: <http://www.tera.com/company/management/smith.html>

Research web page: <http://www.tera.com/company/index.html>

Biographical Sketch:

Burton Smith

Burton Smith is Chief Scientist of Cray Inc. He has held a variety of faculty, federal research, and corporate project leadership posts. He received the BSEE from the University of New Mexico in 1967 and the Sc.D. from MIT in 1972.

From 1985 to 1988 he was Fellow at the Supercomputing Research Center of the Institute for Defense Analyses in Maryland. Before that, he was Vice President of Research and Development at Denelcor, Inc., and was chief architect of the HEP, one of the earliest commercial parallel machines. He developed the Tera Multithreaded Architecture (MTA) and compiler, one of the most innovative parallel architectures ever, with 128 regularly scheduled threads per processor. The company he founded to commercialize the MTA, Tera Computer, bought Cray Inc. on April 27, 2000.

Dr. Smith is a Fellow of both the ACM and the IEEE, and he was winner of the IEEE-ACM Eckert-Mauchly award in 1991. His main interest is general-purpose parallel computer architecture.

He has visited Lawrence Livermore National Laboratory well over a dozen times, beginning in 1979.

April 12, 2000

Partly Random Graphs and Small World Networks

Gilbert Strang

Massachusetts Institute of Technology

Email: gs@math.mit.edu

Abstract:

It is almost true that any two people in the United States are connected by less than six steps from one friend to another. What are models for large graphs with such small diameters? Outstanding applications would be networks of neurons, or an electric power grid, or even (possibly) the world wide web.

Watts and Strogatz observed (in *Nature*, June 1998) that a few random edges in a graph could quickly reduce its diameter (longest distance between two nodes). We report on an analysis by Newman and Watts to estimate the diameter with an N -cycle and M random shortcuts, $1 \ll M \ll N$.

We also study a related model, which adds N edges around a second (but now random) cycle. The average distance between pairs becomes nearly $A \log N + B$. The eigenvalues of the adjacency matrix are surprisingly close to an arithmetic progression; for each cycle they would be cosines, the sum changes the spectrum completely.

We will discuss diameters and eigenvalues of the adjacency matrix for partly random graphs. We also report on the surprising eigenvalue distribution for trees (large and growing) found by Li He and Xiangwei Liu. And a nice work by Jon Kleinberg discusses when the short paths can be located efficiently by a decentralized algorithm—as on the web.

Speaker's web page: <http://www-math.mit.edu/~gs/>

Institution web page: <http://www.mit.edu/>

Biographical Sketch:

Gilbert Strang

Gilbert Strang was an undergraduate at MIT and a Rhodes Scholar at Balliol College, Oxford. His Ph.D. is from UCLA and upon graduation he has taught at MIT. He has been a Sloan Fellow and a Fairchild Scholar and is a Fellow of the American Academy of Arts and Sciences. He is a Professor of Mathematics at MIT and an Honorary Fellow of Balliol College.

Professor Strang has published a monograph with George Fix, "An Analysis of the Finite Element Method," and six textbooks:

- Introduction to Linear Algebra (1993,1998)
- Linear Algebra and Its Applications (1976,1980,1988)
- Introduction to Applied Mathematics (1986)
- Calculus (1991)
- Wavelets and Filter Banks, with Truong Nguyen (1996)
- Linear Algebra, Geodesy, and GPS, with Kai Borre (1997)

He served a two-year term as the President of the Society for Industrial and Applied Mathematics (SIAM), 1999-2000, during which tenure he helped lead SIAM in the formation of an Activity Group in Computational Science & Engineering and organized the first SIAM conference in CS&E.